

**NOAA'S R&D HPCS ACQUISITION
SOLICITATION NUMBER DG1330-05-RP-1038
AMENDMENT 0003
QUESTIONS AND ANSWERS**

Question 29. Section L.6.2.2 states that the Offeror is required to include the following in its cost/price proposal: E. Separate pricing for all of the options described in section C.9. Section L.6.2.5, Funding Profiles, states that Offerors are required to submit cost/price proposals based upon Tables II and III. Table II covers the base contract period, the four year option period, and the one year contract transition period. Table III covers the funding for additional R&D HPCS Augmentations and Engineering Support.

Does the Government want pricing for the following items in C.9 at this time?

C.9.3, One-year extension of Option Period. Since the RFP in Section L.6.2.1 requests pricing for the base period of the contract, does the Government also want pricing for the One-Year Option Period extension at this time?

C.9.4, Additional R&D HPCS Augmentations. This clause states that the Government will request the Contractor to provide a proposal to meet any such requirement. Does the Government want pricing for this option at this time? If yes, what assumptions should be made as to the number and requirements of the potential task order(s) to be issued? Will the Government provide a sample task order to be priced?

C.9.5, Engineering Support. Section L.6.2.5 Table III indicates total pricing per contract year, but Section C.9.5 requests only labor rates. What assumptions should

be made as to requirements of this Option?

Answer: As stated in Section L.6.2.1, Offerors are required to submit price proposals for the Base Contract Period only. Accordingly, Offerors are required to submit price proposals for CLINs 0001, 0002, 0003, 0004, 0005, and 0005A which is an option to extend the Base Contract Period one additional year. In addition, Offerors are required to submit price proposals for CLIN 0010, Additional R&D HPCS Augmentation, and CLIN 0011, Optional Engineering Support Services. Price proposals for both optional CLINs 0010 and 0011 should be for the Base Contract Period only. Price proposals for optional CLINs 0010 and 0011 should be firm-fixed prices for each of the subCLINs identified. For example, what would be the price of one additional Workstream 1 (WS1), Coupled Earth System Model (CLIN 0010A)? What would be the hourly rate for an Applications Analyst (CLIN 0011A)? Also, if discounts are offered for ordering multiple quantities, what would be the discount and the break-point(s)? No sample tasks are being provided for proposal purposes. The dollar figures cited in Section L.6.2.5, Funding Profiles, Table III, expected funding ceilings in \$ millions for the Contract Augmentation and Engineering Support Services are provided for informational purposes and reflect the estimated annual ceiling amount for the optional items during the years specified. No assumptions should be made as to the requirements of these options. Offerors are not required to provide price proposals for the Option Contract Period (CLINs 0006, 0007, 0008, 0009, and 0009A). Nor are offerors required to submit price proposals for the optional items (CLINs 0010 and 0011) during the Option Contract Period.

Question 32. Cover Letter, Page B-5, Section B.2, Note E, Page L-16, section L.6.2.6

In several places (including the RFP cover letter and Note E on page B-5) the Government states that it will procure the HSMS during the first year of the Base Contract Period. It is also stated that archive media will be procured throughout the contract period. Does the Government intend to preclude Contractors from proposing the incremental procurement of archive capacity -- additional silo components -- throughout the period of performance, on an as-needed basis?

Must the Contractor propose the procurement of all archive hardware required during the period of performance during the first year of the contract period? If so, this may require the procurement of components for GFDL and NCEP during the first year of the contract period, when no funds are provided to support these activities.

Alternatively, should we interpret the Government's language to indicate simply that all components of the HSMS (software, hardware, and media) are to be owned by the Government, as opposed to leased?

Answer: The alternative interpretation posed is correct. That is, the Government intends to own, rather than lease, all components of the HSMS (e.g., hardware, software, media, etc.). The Government has deleted mention of acquiring the HSMS during the first years of the Base Contract Period and Option Contract Period from

Section B, Note E, recognizing that upgrades and additional media might be delivered throughout the term of the contract. Offerors are reminded that its price proposals for the Base Contract Period should reflect purchase of the item in the intended year of delivery.

Question 35. Page C-11, Section C.5.2.5

The second paragraph states, the data that is generated from the OCCS, and available at both Fairmont, WV, and Gaithersburg, MD, is required to be written to the R&D HSMS that supports workstreams 4-9. Should this be workstreams 4-6, the workstreams associated with NCEP?

If this is supposed to be workstreams 4-9, are these data currently sent from Fairmont or

Gaithersburg to FSL to support FSL processing? Will these data be sent by OCCS to FSL during the R&D HPCS contract? Is the R&D HPCS contractor responsible for providing that bandwidth? Are these data redundant with the data required to support processing at FSL that are identified on page C-23, Table V - Data Ingest?

Answer: Answer to follow

Question 36. Page C-11, Section C.5.2.6

This section specifies data generation profiles for each site. In order to provide an accurate HSMS sizing, please provide data access profiles for each site (e.g., data volume and number of files retrieved from the HSMS per day).

Answer: NCEP expects an HSMS requirement of at least 2.5 PB by 2007. NCEP staff retrieve many thousands of files each day.

GFDL has calculated the following averages for its HSMS data access profile for the November-Feb time period. There has been an average of 71000 files/day written to and 44000 files/day retrieved from small tapes (9840s). There has been an average of 814GB written to and 670GB retrieved from the small tapes (9840s). There has been an average 7150 files/day written to and 3900 files/day retrieved from large tape (9940s). There has been an average of 9.5TB/day written to and 6.7TB/day read from large tape (9940s).

For FSL, over 200,000 new files are created each year. Over 60TBytes were added in the last year.

Question 61. Found a bug in ws8: wrf_chem - Subroutine SWPARA in

phys/module_ra_sw.F seems to have a bug. It contains the following 2 statements:

```
bexth2o=5.E-6
```

and

```
XSCA=(1.E-5*XATP(K)+aer_dry1(K)*bext340+aer_water1(K)*bexth20)/XMU
```

and they are the only statements involving a variable with "bexth2" in its name. It would seem the "bexth20" in the second statement has a typo ("zero" instead of "o"), so that aer_water1(K) will get multiplied by whatever trash happens to be in bexth20 (there is no expectation that it will contain 0 or some other fixed value, as it would be allocated from the stack and in general not saved from call to call). Please confirm whether you acknowledge this as a bug. An estimate of what the impact of this could be on the results would be appreciated.

Answer: As you correctly determined the variable bexth20 should have been bexth2o (with an oh, not a zero). This change has very little impact on the results, as bexth2o is small (5E-6). The corrected routine 'module_ra_sw.F', as well as the output from NetCDFcompare for T and P (temperature and pressure) are posted on the Q&A webpage. These variables were not effected by the change.

Question 63. I did verification for one output and it shows "0 variable processed"

and "passed". Does that mean everything is ok?

Answer: New verification procedures, using layered averages, will be posted on 3/4/5.

Question 87. Section C.10.3.2 provides current archive data holding sizes for each site and Section C.12 provides the list of GFE available at each site. As indicated in Section

C.10.3.2.3 Figure 3, data holdings will continue to grow at GFDL, and presumably at other Sites, during the period leading up to the initial system delivery under this contract.

Presumably some archive configurations will need to be enhanced to accommodate this growth. Please provide guidance on what assumptions we should make about the size of data holdings and the archive configurations that will exist at each site at the time of initial delivery.

Answer: Answer to follow

Question 94. Facility Power Service Clarification - Section H.19 and Section C.11.3.1 Section H.19 states, "The Government will provide the Contractor with credits for monthly power usage, as indicated by the total kilowatt hours in the PSE &G bill, that exceed two times the monthly power usage for the Computer Building."

Section C.11.3.1 indicates that the underground PSE&G 4160-volt feeder dedicated to the Princeton Complex is routed to separate building substations (each located in the Princeton Main Building and the Computer Building). The Government indicates that Main Building substation provides power to the Main building and the Chilled Water Plant and the second substation to the Computer building. The Government states that the Main building Substation utilization is dominated by the Chilled Water Plants. The Bidder requires the following clarification:

1) In Section H.19, does the Government's statement that it will provide Contractor credits for monthly power usage in excess of two times the monthly power usage for the Computer Building assume that the power consumed by the Main Building substation is roughly equivalent to power consumption usage of the Computer Building substation?

2) If this statement is true, then should the Contractor have to upgrade the Chilled Water

Plant located in the Main Building, would the total power usage the Contractor was responsible to pay for still be equivalent of 2x that measured by the Computer Building power meter? In other words, are the Power Cost the contractor is responsible for figured as a factor of 2 times the metered usage of the Computer Building Substation rather than the actual metered usage of both the Main Building and the Computer Building substations?

3) Section C.11.3.2 provides a table containing "Total Usage and Expenditures for Electrical Utilities of the Princeton Complex." In C.11.3.1 the Princeton Complex is

defined as containing a substation providing power support to the Princeton Complex as well as to several other buildings on the B site of the Forrestal Campus. The Government further states, "Power usage for these other buildings is primarily offices, although research activities in some of these other buildings occasionally require substantial power from the substation." Are the Figures provided in the table referenced in this question inclusive of "other buildings on the B site or the Forrestal Campus" and do they include any "research Activities" unrelated to the requirements of the Chiller Plant and Computer Building? If so, is it required that the Contractor include these additional costs as a part of their price proposal?

4) If the Government elects to retain existing, currently installed GFE and Vendor or Integrator owned equipment beyond the start of the NOAA R & D contract, will the Contractor be responsible for paying for the electrical utility costs associated with this equipment beyond the start of the NOAA R & D contract. If so, how will these costs be determined?

5) Please clarify if the electrical utility costs for BLDR-1 and BLDR-2 are the responsibility of the Contractor. Since the costs provided in the RFP by the Government for BLDR-1 and BLDR-2 are only estimates, and the Government states that should metering occur sometime in the future, costs could increase substantially, how can the Bidder be asked to provide a firm, fixed price for this service when the Government cannot define the potential or actual cost. This appears to place an undue risk upon the Contractor.

Answer: 1.) Based on Amendment 2, there is no assumption regarding the power

consumption by the two substations, other than that the total power consumption of the Princeton Complex will be greater than two times that of the equipment operated under the R&D contract.

2.) See Amendment 2.

3.) No, the power consumption of the rest of the B-site campus is on a separate meter that is the responsibility of Princeton University.

4.) If he uses the PRTN facility, the Contractor shall only be responsible for paying the electrical utility costs for the operation and associated cooling of equipment operating at the facility under the R&D contract.

5) The electrical utility costs for BLDR-1 and BLDR-2 are the responsibility of the Government. The per KWH costs were provided for comparison only, should the bidder propose a site other than BLDR-1 or BLDR-2. No savings recouped by not using BLDR-1 or BLDR-2 will be added to the contract.

Question 100. Reference C.10.3.2 Hierarchical Storage Management System (HSMS) -

NOAA's FSL site currently uses LTO technology while GFDL and NCEP use StorageTek 9940 and 9840 tape technology. In the publication "The IBM TotalStorage Tape Selection and Differentiation Guide" (www.ibm.com/redbooks) page 38, this manufacturer of LTO technology states that LTO will work but is not recommended as a medium for HSM applications.

Will the Government accept LTO for the high performance HSMS system?

Is it desirable to have the same tape technology at all three sites?

Answer: Section C.5.2.4 states that the offline tier is robotically mounted but infrequently accessed. Enterprise Backup class robots, drives (such as LTO), and media may be suitable for this application.

The workstream benchmark, data generation, and data retention profiles are meant to give an idea of performance and capability requirements by the Government for the HSMS. These metrics are given on a workstream by workstream basis so that an Offerer may provide a solution that balances all storage needs for the given workstream.

NOAA's current HPC centers use a mixture of drive types. And as appropriate to the capacity, performance, and duty cycle requirements, LTO has been used. See Section C.10.3.2 of the rfp for more details.

The Government does not require the same technology to be used to solve each workstreams HSMS needs.

Question 105. Re: WS8 & WS9 verification question - Following your response to question 49, we modified the stats.inc file as indicated. Unfortunately, we still cannot verify our output for either WS8 or WS9. For both cases we encounter grid point values that do not pass verification, although the global averages are otherwise quite close. Looking at the case for WS8 we see values such as the following in the NetCDFcompare output (just picking out levels with maximum normalized differences):

Opening file: wrfout_d01_001200_ref

Opening file: wrfout_d01_001200_64p

155 variables processed

(skipping lots of output)

P: Mean A = 1.00680E+03, Mean B = 1.00642E+03, Norm Diff = 3.76449E-04

Level	MAX	MIN	MEAN	STD DEV
-------	-----	-----	------	---------

MAX DIFF	MAX DIFF MAG	MEAN DIFF	STD DEV DIFF
----------	--------------	-----------	--------------

NORMAL DIFF

1	2.34510E+03	9.42037E+02	1.62073E+03	2.93297E+02
---	-------------	-------------	-------------	-------------

	4.03270E+02	2.34510E+03	5.62318E+00	7.62088E+00	1.71963E-01
--	-------------	-------------	-------------	-------------	-------------

```

2  2.32498E+03  9.36850E+02  1.61304E+03  2.92501E+02
    8.60933E+01  1.61181E+03  5.21517E+00  6.10921E+00  5.34139E-02
.
.
.

10 2.11268E+03  7.90449E+02  1.43258E+03  2.79140E+02
    1.34863E+02  1.11574E+03  6.00967E+00  9.25817E+00  1.20873E-01

11 2.06996E+03  7.67445E+02  1.39934E+03  2.76605E+02
    2.33016E+02  1.34117E+03  6.36916E+00  1.20507E+01  1.73740E-01

12 2.04214E+03  7.52130E+02  1.37729E+03  2.74477E+02
    2.79264E+02  1.75708E+03  5.47399E+00  9.61002E+00  1.58937E-01
.
.
.

```

(skipping lots of output)

T: Mean A = 1.78069E+01, Mean B = 1.78305E+01, Norm Diff = -1.32429E-03

Level	MAX	MIN	MEAN	STD DEV	
	MAX DIFF	MAX DIFF MAG	MEAN DIFF	STD DEV DIFF	
	NORMAL DIFF				
1	1.83693E+01	-1.76940E+01	-2.16531E+00	6.57712E+00	
	5.28220E+00	1.14892E+01	2.08118E-01	3.29907E-01	4.59755E-01
.					
.					
.					
8	1.78622E+01	-1.54519E+01	-2.15474E+00	6.36155E+00	
	2.24063E+00	1.76791E+00	1.50625E-01	2.02123E-01	1.26739E+00
9	1.79109E+01	-1.51969E+01	-1.79414E+00	6.38714E+00	
	2.10944E+00	2.32166E+00	1.40706E-01	1.85165E-01	9.08591E-01
10	1.80265E+01	-1.49247E+01	-1.28641E+00	6.39764E+00	
	2.11823E+00	2.44168E+00	1.15795E-01	1.55444E-01	8.67527E-01
.					
.					

The situation with WS9 is similar. Reduction of optimization only changed results marginally. Can you help us resolve this issue?

Answer: New verification procedures, using layer averages, for WS8 and WS9 were posted on 3/4/5.

Question 107. Section C.5.5 (Page 24) states the agency must provide a security plan for each General Support System. Section H.16(c) states for all Contractor-owned systems, the Contractor must provide an IT Security Plan and Section H.16 (c) 2 states the Contractor must submit a System Certification and Accreditation package.

In the context of the security requirements and responsibilities as referenced above, does the Government consider the leased systems contemplated for this procurement as Contractor-owned or Government-owned?

Answer: In the Amendment 0002, Section C.5.5 has been updated:

All IT equipment delivered as part of this procurement is to be considered a Government computing resource, regardless of its location or actual owner. The Government will initiate the C&A package for this system and then transition

maintenance of the security documentation to the Contractor after delivery and acceptance. The Contractor may also elect to employ equipment which is not part of the delivery for use in remotely diagnosing, monitoring, or managing the delivered system. If this Contractor-owned equipment has any special access or trust relationship with the delivered system, then the Contractor must initiate and maintain any appropriate C&A documentation for this equipment as a separate Government System.

Question 112. In C.5.2.4, p. 8, requirements are stated to keep the existing data in the HSMS archives available. It may be necessary to re-archive this data on new media in order to maintain it. Please provide estimates for the following:

- a) The projected size of the HSMS data archive in Princeton at the start of FY2007 and the retention requirements for this data;
- b) The projected size of the HSMS data archive in Gaithersburg at the start of FY2007 and the retention requirements for this data;
- c) The projected size of the HSMS data archive in Boulder at the start of FY2006 and the retention requirements for this data.

Answer: NCEP plans to limit its future storage and not add additional silos its storage system. NCEP plans to turn over two full silos from Gaithersburg, some 2.5 PB of data, to the R&D HPCS in FY 2007. One additional silo (1.25 PB) in Fairmont

will also be turned over. Data retention policies are described in Section C. NCEP estimates that a total of 2.5 PB of data will be retained by the R&D HPCS (see Q&A #43).

At the Boulder site, the HSMS will be maintained by the current vendor at the Government's cost through the first year of the new contract. The data retention policies are described in section C.

At the Princeton site, as mentioned in question 87, it is anticipated that by the start of FY2007 there will be slightly more than 7PB of data in 5 StorageTek silos using 9840, 9940 and Titanium media. All of this data will need to be retained for at least nine years.

Question 113. WS 8 verification - After changing tolerance to 0.1 as noted in answer to question 49 NetCDFcompare still fails to pass for fields T and P in WS8. The results do appear to be well within 10-20% requirement however. Can you check the results listed below for temperature and verify?

I cannot post results in this web form - I get errors. Where can I send NetCDFcompare results for inspection?

Answer: New verification procedures, using layer averages of P and T, will be posted

on 3/4/05.

Question 118. Section C.8.4 Configuration and Change Management Plan states that, THE GOVERNMENT DESIRES THAT THE CONTROL BOARD AND CONFIGURATION MANAGEMENT PROCESS INTEGRATE INTO EXISTING GOVERNMENT CONTROL BOARDS AND PROCESSES WHERE APPLICABLE. Please provide examples of processes, procedures, and Control Board and Configuration Management functions currently performed at each R&D laboratory today.

Answer: Answer to follow

Question 125. In Section C.11.4.1, for the Princeton site, the Government expresses some doubt as to the ability to operate all the Chillers concurrently without modifying the flow sensitivity controls. It also mentions a potential redesign of the chilled water plumbing and/or controls and states that until this situation is resolved, operation appears to be limited to chillers #3 and #4 for heating loads requiring more than one chiller. Would the Government please clarify exactly how much chiller capability will be available to the Contractor at the time of initial installation? Furthermore, would the government please clarify who will be responsible for the redesign of the chilled water plumbing and/or controls and the associated cost of that redesign and implementation? Additionally, does the Government have a minimum set of

standards relative to the design and build-out for any increased chiller plant capacity and what is the Governments minimum requirement for backup capacity? Does the Government have any NOAA-wide standard required to Backup and Recovery standards as relates to Power and Cooling and generator capacity? Should the Contractor propose an alternative Contractor-provided facility does some set of minimum standards regarding Section 11 capabilities exist?

Answer: Answer to follow

Question 139. Scaling Study Requirements - Can the Government clarify the requirements with respect to the data points gathered for the scaling study?

Answer: The goal of the PE count elements A. - I. in section J.1.4.3.2 is to require a scaling study data point at the PE count proposed for the throughput test. However, the phrasing of the requirement assumes that the part of the system offered for a given workstream has all instances of the workstream throughput running concurrently as the basis for the proposed throughput performance.

On the other hand, there is nothing which actually requires concurrence of all workstream instances. The requirement is simply to provide the best throughput for the resources associated with a given workstream.

To simplify the statement and make it sufficiently general, the PE count requirement as expressed in items A. - I. of section J.1.4.3.2 will be amended to read:

1) One of the scaling study points must be provided at the PE count proposed for the throughput. In the event that multiple throughput PE counts are proposed, the Offeror may use scaling study data from multiple throughput PE counts to fulfill the required minimum number of scaling data points if the proposed throughput PE counts differ by 50%. It is of no use to the Government to have "scaling data" with data points taken at roughly similar PE counts.

2) In the event that all instances of a given throughput workstream do not run concurrently on the proposed system, the Offeror must provide either:

a) a scaling data point for the PE count which would allow the workstream instances to run concurrently

b) a statement describing the technical reasons why the resources provided for a given workstream cannot support all instances of that workstream.

Questions 142 and 143. For WS8, let us assume that the only test case that the

Government is asking for is the em_real test case but note that question 58 has not been answered. Verification remains problematic because the numerical computation for WS8 is demonstrably unreliable with 4-byte real variables. We have explored this by enabling or disabling combined multiply-add operations. The computation without combined multiply-add operations results in IEEE-compliant single-precision results, and enabling combined multiply-add operations results in numerical precision that exceeds the IEEE specification. Comparison of these two sets of results from the same machine shows that the predicted temperatures can vary by several degrees when such slight changes in arithmetic are made. For the temperature field, the average difference remains fairly small (in units of Kelvin) for each level, but there are grid points where the temperature difference is substantial. Similar differences arise with other variables. According to the benchmark documents, the reference data comes from a system that is not IEEE compliant and further, numerical results on X86 systems vary depending on which floating-point hardware (X87 or SSE) operates on the data. Given the demonstrated sensitivity of the numerical results to the details of floating-point arithmetic, a different verification process is required. Something similar to the solution proposed in Question 133 seems reasonable, but given the observed numerical sensitivity, it may be difficult to define a meaningful verification method. We have no doubt that the application is ported correctly, but when the results are this sensitive, how can they be verified?

Answer: Only em_real is required for WS8. Test cases are provided to ensure correct execution. New verification procedures, using layered averages, will be

posted on 3/04/05.

Question 146. Please clarify the data retention period for workstreams WS1, WS2, and WS3 as listed in Table III of Section C.5.2.7. The table specifies that the workstream data is to be 100% retained for the life of the NOAA R & D contract. It further states that 50% of this is to be retained as a persistent archive. Please clarify the difference between 100% retained, and 50% in persistent archive. These appear to be mutually exclusive directives.

Answer: 100% of the data must be retained throughout the life of the contract.

"Persistent archive" will be added to the C.13.2 definitions in an amendment.

Question 150. Section E.2.3.2 Disk I/O Performance, specifies that "When possible, all performance specifications indicated in the contractor's proposal will be verified. Any failures to meet the specifications shall be remedied or result in failure of the acceptance test." There are no disk performance benchmarks provided as part of this RFP. Will

the disk I/O performance be measured solely against the vendor's provided specifications? Should this section of the RFP be modified or removed?

Answer: Answer to follow

Question 151. The new RUC_WS data provided is for different date than the original file, so bidders will need a new file to verify model output against. Will that be provided? (We believe it should be 050381248.NNT_dat). Additionally, the new data release does not sufficiently answer the questions posed in submitted questions #46 & 47. Will those questions be answered directly?

Answer: The file which should be verified is the post-processed 050381248.NNT_dat file. The post processor will produce a maps_fcst/0503812000048.grib file which can be compared to the maps_fcst/0503812000048.grib.save file provided. A number of grib utilities are available online. We use wgrib from NCEP website which allows you to dump grib headers and variables as ASCII or binary formatted files.

<http://www.cpc.ncep.noaa.gov/products/wesley/wgrib.html>

A file "RUC 20 file information" was posted to the NOAA HPCS Benchmark information page on 2/23/05 which addresses question 46 and 47.

Question 158. Fairmont physical facility upgrades - To what extent will the contractor be responsible for physical plant upgrades at Fairmont.

Answer: The Government will provide 450 KVA of power and 130 tons of cooling to support an R & D system in Fairmont, at no cost to the vendor. Incremental facility resources greater than those amounts will be the responsibility of the vendor. These specifications take into account the needs of the Operational System.

Question 159. Data communications - Please describe where the workstreams can be geographically located without incurring any communications costs.

Answer: WS 1-3: Princeton.

WS 4-6: DC metro and Fairmont.

WS 7-9: Boulder.

For WS 4-6, a fraction (approximately 50%) of the data exists in Boulder.

For WS 7-9: a fraction (approximately 80%) of the data exists in DC metro and Fairmont.

Question 160. Will NOAA consider performance results that are obtained from images built with profile guided optimization as CLASS A results?

To provide a bit more detail,

- Profile Guided Optimization does not involve any source code modifications

- Profile Guided optimization can materially improve the performance of the executable

- Basically, the process is that you build the application TWICE, once with a base set of optimizations, then you run one or more sample executions, then you rebuild the application based upon the execution profile. The new image has optimizations that can be substantially better than the original image.

- Profile Guided Optimization is a documented, supported feature in our compilers (and many alternative compilers as well). Some have different names for this facility, but they are quite similar in process to utilize and impact of results.

Answer: The Government encourages the use of innovative technologies to produce better code optimization. On the other hand, the point of the performance baseline based on class A changes alone is to help the Government understand the cost/benefit of anything more than basic compilation.

Profile Guided Optimization is a technology long present in a number of compilers.

But to date, the Government has had little or no success on real applications with previous implementations. Moreover, since various features of the model code are activated depending on events occurring within the simulation run, the question arises as to how much data must be sampled to arrive at an overall improvement over the

entire simulation length. Further, how generally applicable are the enhanced compiler settings? Must one run and re-run portions of new simulations to get useful compiler settings? How long must one run to obtain sufficient performance data? In net, how much human resource will be required to keep the results of profile guided optimization up to date and useful?

Thus in order to understand the benefit, the Government views the use of profile guided optimization as a class C change. Further the documentation process required for class C changes is designed to help the Government understand the costs of implementation as well as give some insight into the generality of the technology. Thus should an Offeror choose to provide data using profile guided optimization, the Offeror must address the issues described above in their proposal.

Note that there is no inherently negative connotation to either class B or class C changes. The Government simply needs to assess the cost/benefit

Question 162. Section C.5.1.2, Development Component - If the resources provided for the Development Component are identical to the resources provided for the Large-Scale Computing Component, and the Resource Management Software allows these resources to be added to or subtracted from the batch pool dynamically and interactively, can the Development Component resources be counted in the processor count when calculating/extrapolating the total workstream throughput times for the

proposed system -

or do all resources proposed for the Development Component have to be ignored when calculating the LSC system throughput?

Answer: The Offeror must be specific about the number of processors added to the batch pool, the performance increase attributed to them and must provide a complete description of the process required to include and remove them (return to other tasks) from batch processing. If these conditions are met, the Offeror is allowed to include these processors in LSC throughput calculations.

Question 163. In section C.5.2.6.1 the Government provides data generation rates for workstreams 1 through 3 (1.4 TB/day, 2.6 TB/day, and 1 TB/day, respectively) for a baseline level of performance for these workstreams. Do these data generation rates correspond to the current Origin LSC baseline or to the Altix LSC baseline to be installed in April 2005?

If these data generation rates correspond to the Origin baseline, for the purpose of projecting data generation into the future, should we assume that the Origin baseline has a relative performance level of 1.0, and the Altix baseline has a relative performance level of 1.8, as stated in appendix A (C.10.1.3, page 40)? Such an assumption is necessary because the Government did not provide throughput benchmarks (section J.1.4.2.4) for the Origin baseline and therefore our performance

baseline comparisons can only be made to the Altix baseline.

If this assumption is correct, then the data generation rates expected from the Altix system beginning in April 2005 would be 1.5 (1.8 raised to the 0.7 power) times the values stated above, or 2.1 TB/day, 3.9 TB/day, and 1.5 TB/day respectively -- correct?

Answer: Answer to follow

Question 169. Regarding the verification for WS6, GSI, the README file says,

At the completion of the analysis, we would like to see 3-4 digits accuracy in surface pressure and in the rms of the temperature fields when comparing column 1 with column 2. The third column, rms of the differences of the temperature fields, should be less than 0.5

1) There is only reference to the rms values of temperature fields. Are the other fields subject to the same verification criteria?

2) If the above criteria apply to the other fields, the first point of matching 3-4 digits of

accuracy between columns one and two makes sense. However, the way rms.diff.f

computes the third column is to compute the rms value of the difference between columns 1 and 2. Unnormalized, it is possible for the vendor solution to be off by several orders of magnitude (violating the first criterion), but have the rms values in the third column be less than 0.5. Say if the verification rms value is $10E-7$, the vendor rms value is $10E-4$, the difference will be roughly $10E-4$ which is less than 0.5. Can you clarify this, or specifically state the verification criteria for the fields other than the temperature?

Answer: No, the other fields are not subject to verification.

Question 170. When will the 7052 sq ft of floor space in Princeton be available?

Who will be responsible to remove existing SGI equipment? Will the selected vendor have access to the facility prior to removal of the SGI equipment to perform any facilities upgrades necessary to support the new hardware platform?

Answer: Answer to follow

Question 171. Table III in paragraph C.5.2.7 provides the data retention profile for all

workstreams. The paragraph and table use the term "persistent archive."

1) Please provide your definition of the term "persistent archive."

2) Please define the requirements for the persistent archive in terms of permissible location, recall requirements, and performance.

Answer: 1) "Persistent Archive" is the data in the archive that will exist or remain indefinitely. 2) The data retention profile does not imply use characteristics. Please see section C.5.2.4 for HSM requirements.

Question 172 and 173. Could the government please describe the 3296 processor Altix configuration scheduled for delivery to the Princeton site in April 2005? Will this be a single system image or will there be several clusters? If there are several clusters , what size will each cluster be and will they be connected via a high speed interconnect? If it is several clusters , does the government expect to run jobs across the clusters?

Answer: The April 2005 configuration for the Princeton site LSC will be comprised of 512, 256 and 96 processor SGI Altix Systems. Thus, the configuration is a cluster of systems and there is no single system image. Further, the Government currently has no plans to run applications outside a single system's numalink communication fabric.

Question 174. In an effort to optimize networking costs for the Government, can the Government describe any existing, or near-term planned, fiber network connection points (speed and type) that may be in close proximity (less than 2 or 3 km) to the current GFDL computing facility? Are there any fiber connection points, such as for Dept of Energy or Physics Departments, that could be accessible and with available bandwidth?

Answer: There is a 36-strand fiber bundle between NOAA's GFDL facility and Princeton University's Sayre Hall. It contains 9 single-mode pairs and 9 multi-mode pairs.

There is a 12 strand fiber cable between Princeton's Sayre Hall and DOE's Princeton Plasma Physics Laboratory (PPPL). One pair currently supports the Microwave connection at Sayre Hall and 2 pairs are reserved, leaving 3 pair dark. There is an easement protecting the fiber route path which extends roughly 1 mile between Sayre Hall and PPPL.

Question 175. The answers to several questions refer to an amended version of the RFP. When will this be available? It would be very helpful to have an amended version

that integrates the answers to all questions posted so far as soon as possible.

Answer: Amendment 0002 is currently being prepared and should be released by March 15, 2005. Amendment 0002 incorporates several changes to Sections C and J resulting from responses to vendors' questions. Please monitor the following URL for the Amendment: <http://www.rdc.noaa.gov/~amd/index.html> Click on "Solicitations" and scroll down to "NOAA's R&D HPCS."

Question 176. Describe the access route for moving equipment from the loading dock to BLDR-2.

Answer: Answer to follow

Question 177. Describe the power panels in BLDR-1 and the way in which the power is distributed under the raised floor.

Answer: In BLDR-1, power is fed from the switchgear to UPS systems located within the room. These systems support 480V in and 480V out. Power then feeds to a transformer, which feeds a wall mounted power panel. The power panels then feed underfloor junction boxes with 50ft. flexible conduit receptacles of various configurations.

Question 178. What are the door clearances along the access paths from the loading

dock to BLDR-1 and BLDR-2?

Answer: See answer to question #176.

Question 179. What are the issues that will need to be addressed in order to install an additional chiller in the chiller room (Boulder)?

Answer: Answer to follow

Question 180. What power is available to vendors in BLDR-1 and BLDR-2?

Answer: Power in BLDR-1 (available October 2006) will be 250kVA. Power in BLDR-2 (available October 2005) will be 350kVA.

Please refer to Section C.11.11

Question 181. Why is the right side of the raised floor in BLDR-2 not available for this contract?

Answer: That area is designated for non-HPC equipment that will support other NOAA projects.

Question 182. Will the UPS systems to be installed in BLDR-2 use chloride?

Answer: The construction contract for the BLDR-2 site has not been awarded, and therefore the brand of UPS is unknown.

Question 183. Who is responsible for the maintenance of UPS and CRAC units (both GFE and new) that are used in this contract?

Answer: All UPS systems are maintained by the government. All CRAC units chill water systems are maintained by the GSA building maintenance contractor on site.

Question 184. What options are available for changing the positions of the XDO units to be installed in BLDR-2? If a vendor chooses not to use the XDO units, what ceiling clearances will be allowed for the systems that he/she installs?

Answer: Amendment #3 will make figure #2A available which shows that the XDO units have been relocated. They are now designed into a position that will support racks up to 48" deep, while still allowing for 48" separation in the "cold aisle".

Equipment not located under the XDO units will not receive any overhead cooling, and therefore should adhere to the 84" restriction. Additional guidance is provided in Section C.11.1.8

Question 185. Does the Government have a plan for implementing dual power configurations?

Answer: The Boulder sites do not have the necessary power distribution to support dual power configurations as a form of redundancy.

Question 186. How many engineers does the current HPC vendor in Boulder have on-site? Does the current HPC vendor have any other support personnel for the current system?

Answer: Answer to follow

Question 187. I saw a rack in the BLDR-1 computer room that had a ceiling clearance of only about 4 inches. Why does this rack not conform to your stated specs?

Answer: BLDR-1 is in the process of phasing out all of the racks that do not meet the required specifications. The rack you saw is one of the few remaining to be removed and or replaced.

Question 188. Will the BLDR-2 fire suppression system be tied into the corresponding BLDR-1 system?

Answer: The BLDR-2 system will not be tied to the BLDR-1 system. However, both systems are tied to the DSRC Building Automation System.

Question 189. Where will the Telco demarks be located for BLDR-1 and BLDR-2?

Answer: The telco demark is the Main Distribution Facility (MDF), but often times the service provider will extend the demark to the computer rooms.

Question 190. Describe the process by which a vendor will implement any facility build-outs that he/she requires. Who will pay for these build-outs and how will payments be made? Can the Government provide guidance on the estimated time (in weeks or months) required to complete build-outs of different degrees of complexity?

Answer: Answer to follow

Question 191. What is the cost of a fiber connection from the MDF to BLDR-2?

Answer: Answer to follow

Question 192. What LAN connections are available in BLDR-1? 100Mbps?
Gigabit?

Answer: Answer to follow

Question 193. When the labs are combined into ESRL, will there be a change in
lab/production status?

Answer: Answer to follow

Question 194. Does the BLDR-2 fiber run from MDF?

Answer: Yes, the BLDR-2 fiber will run from the Main Distribution Facility (MDF)
in accordance with building standards.

Question 195. What mix of vendor engineers are on-site?

Answer: Answer to follow

Question 196. Is the NLR connection to be for BLDR-1 and BLDR-2?

Answer: Answer to follow

Question 197. How long must the newly upgraded equipment in BLDR-1 coexist with any new equipment (installed under this new contract)?

Answer: Answer to follow

Question 198. In the case of a chiller upgrade, who bears the cost? Is this the vendor or NOAA?

Answer: In the case of a chiller upgrade, in the context of this contract, the cost would be born by the vendor in the form of credits or holdbacks. NOAA does not currently have any plans to upgrade the chiller plant at this time.

Question 199. For the GSA build-out of BLDR-2, where will the electrical (wiring) terminate? Will it stop at the panel or will it include electrical whips under the floor?

Answer: BLDR-2 is designed to have power distribution under the floor for each rack. See question #137 for the receptacle configuration.

Question 200. Does GSA have a construction contractor on board to do the BLDR-2 work or is the construction to be bid out when construction documents are finalized?

Answer: Once construction documents are finalized, GSA will hire a construction contractor in accordance with normal government contracting guidelines.

Question 201. Will GSA conduct design and construction management/bidding for the

Vendor's upgrades (under this contract)? If so, will the Vendor be a part of the plan review/approval process and the evaluation of construction bids?

Answer: Answer to follow.

Question 204. Is there a strategy for accommodating dual power and DP (dual power) equipment? For example, two independent power paths from redundant

UPSs? Or independent alternate power sources (say, from the utility)?

Answer: There is currently no strategy for accommodating dual power equipment from an increased reliability stand point. As a research center, there is no requirement for redundant UPS systems or independent utility power sources.

Question 205. (Boulder facility) What is the organization/staffing and the function of the 12-hour (7am - 7pm) support organization mentioned during the tour?

Answer: The organization mentioned is our Systems Support Group. Their function and staffing is provided as part of their web site: <http://www-its.fsl.noaa.gov/ssg/>

Question 206. Have the 13 four-ton XDO systems been de-rated for altitude?

Answer: Answer to follow

Question 207. Where will the power distribution panels be located in BLDR-2?

Answer: The design of BLDR-2 will not be finalized until mid-April. Power distribution panel locations will be known at that time.

Question 208. Is the power/cooling equipment in BLDR-2 fixed, or can suggestions be made on (where) to relocate (it)?

Answer: The power and cooling equipment in BLDR-2 has been positioned to accommodate the best use of floor space, allowing for proper clearances, cost implications and design requirements.

Question 209. Regarding the ramp to BLDR-2, how will it be constructed and what will its load-rating be?

Answer: This ramp section is still under design. However, the material for this ramp will be concrete, but the slope will be 1:12 in accordance with the Americans with Disabilities Act.

Question 210. (Boulder facility) Will the mechanical, electrical, and network drawings be made available?

Answer: Answer to follow

Question 211. (Boulder facility) If upgrades to power and cooling are required,

where will they be located?

Answer: Answer to follow

Question 212. In the answer to Question 69 (answered by NOAA on 2/11/05: It is stated the commensurate cooling is 72 tons in BLDR-1, along with the offered UPS power of 250 kVA. This seems low considering the CRAC units (4) installed in the room. Is this a limitation of the Building chilled water system?

Answer: No. The installed CRAC units in BLDR-1 are currently cooling 300kVA of equipment. Only 250kVA of power is available for HPC, while the other 50kVA power is consumed by other laboratory interests. Therefore, the (derated) cooling available is 72 tons for 250kVA.

Question 213. Can the location of the ceiling XDO cooling units in BLDR-2 be moved 2 feet or 4 feet toward the middle of the room? (i.e., toward the columns?)

Answer: Figure 2A in Amendment #3 shows that the design of the XDO units has them moved toward the columns in the middle of the room.

Question 214. (Boulder facility) It appears that XDO-16 units (four-ton) are

specified. Can eight-ton XDO units be installed?

Answer: The design currently specifies the XDO-16 units and the government feels that this will be appropriate for the amount of power that is offered.

Question 215. (Boulder facility) Is the lowest door height (along the access path) at least 80 inches? This is for access from the loading dock to the inside of both computer rooms, BLDR-1 and BLDR-2. If the door height is less than 80 inches, can the air seal be temporarily removed to increase clearance?

Answer: The door frame height, at the lowest point along the pathway to either BLDR-1 or BLDR-2 is 83". Please see the answer to question #176 for further information.

Question 216. For sub-floor branch circuit conduit in BLDR-1, can computer grade blue flex conduit (Ultratight) which terminates on the circuit breaker panel enclosure be utilized (rather than the current arrangement of Greenfield terminating in a j-box)?

Answer: Answer to follow

Question 217. In section J.1.4.2.1 the government stress that "Multiple workstream

suites with a shared IT architecture target must run together in order to achieve the performance level proposed of each workstream suite. Contractors are cautioned that additional system components will be required during the contract should a workstream fail to meet the proposed performance." It appears that the provided government throughput timings ran only a single instance of the throughput workstreams rather than multiple instances. The workstream timings thus do not take into effect the interaction of the workstreams with each other. In particular, it is difficult to understand how the government will efficiently use the 3296 Altix resource to run workstreams 1-3 using the throughput processor counts detailed in J.1.4.2.4 unless the system is a single system with shared resources. Can the government please clarify this for each site?

Answer: It is the case that but a single instance of each throughput workstream was run to provide baseline performance values for Section J. But it must be remembered that the benchmarks are but surrogates for a much larger body of applications. Moreover, the test cases were run on systems running their normal production work. Thus, the Government fails to see the relevance that the particular background load contending for resources with the benchmark test cases was not another instance of a benchmark test case.

The timings provided with Section J are intended to illustrate current application performance and provide a baseline against which proposed performance can be evaluated. Moreover, job granularity limitations of the current systems are not

relevant to the goals for the systems proposed and delivered for this RFP.

Question 228. Will NOAA reconsider the 12-point Times Roman font requirement for Figures and diagrams? We understand that NOAA requires all text to be in 12-point. But requiring 12-point for labels in Figures and diagrams greatly limits the ability to identify data in a meaningful and clear way.

Answer: Diagrams and labels within figures need not be submitted in 12 point font. However, the font size used must be readable.

Question 231. Section C.11.3.3 identifies two UPS units at PRTN - a 225 kVA UPS owned by the Government and a 500 kVA unit owned by the incumbent contractor. Please clarify the Government's intentions regarding the disposition of the incumbent-owned UPS if the incumbent is not the successful offeror on this solicitation.

Answer: At this time, the Government has no plans for the 500kVA UPS. Offerers interested in the UPS should contact Raytheon directly.

Question 238. Please define WS4 WRF-NMM throughput baseline from section J.1.4.2.4. Is an instance one copy of wrf.exe? What is meant by 3 sequential runs? Is it 3 copies of wrf.exe followed by 3 more then followed by 2 copies? Or is it 24

copies, 3x8, followed by 24 more, etc.?

Answer: WS4 throughput benchmark is defined as follows:

An instance consists of three sequential runs of wrf.exe with forecast length of 48 hours. In other words, run wrf.exe - upon completion run wrf.exe again (2nd)- upon completion run wrf.exe again (3rd).

WS4 requires 8 instances to run.

In Section J.1.4.2.4, the baseline total throughput time is 20478 seconds.

Question 239. Amendment 002 changes C.5.2.4 to allow disk storage for the nearline tier. However J.1.4.4, the HSMS Archive Benchmark, assumes that the nearline tier will be tape storage. Will J.1.4.4 be updated to explicitly cover a nearline tier that is disk storage?

Answer: An amendment will update the HSMS Archive Benchmark (J.1.4.4) to reflect the changes in Amendment 002 C.5.2.4.

Question 240. RE: HIMF scaling study - The benchmark instructions for the scaling study on page 19 of section J say the government requires that at least one of the data

points for workstream 3 be run on 1/2 of the model application PEs proposed as the target architecture for that workstream. Can you please clarify whether that would mean running on $3296/2=1648$ processors or $512/2=256$ processors if the target architecture was the system described in answer 172?

Answer: The requirement for all workstreams including WS3 has been amended for simplification. Please see Benchmark Q&A 139.

Question 241. What is the criterion for answering the "Reproduces across PEs?" question in the WS7 Benchmark Performance Results table? Are bitwise identical outputs expected or required? Is it intended that the bitwise exact reduction mode of SMS be used (even though it was not used in the example scripts or outputs)?

Answer: When compiled with -O0 and with bitwise exact SMS switch set it should produce binary reproducibility across various configurations. However for the scaling study this is not required, and more aggressive optimization is encouraged. The verification criteria, as provided with the benchmark, is the only requirement.

Question 242. There seems to be a discrepancy between sections C.5.2.4 and J.1.4.4.1. Section C.5.2.4(Amendment 2) now allows the use of disk storage for the nearline tier (component of) as long as the data is recoverable via 2nd copy/backups. Section J.1.4.4.1 (WS1-3) benchmark still requires that archival files be restored from

tape media. These sections appear to be at odds with one another.

Answer: Please see the Benchmark RFP Question 239.

Question 243. The RFP discusses (J.1.4.2.1) running more than one workstreams concurrently for throughput tests. Could you clarify whether vendors are required to run more than one workstream (even groups such as WS1-3) concurrently as part of the throughput tests?

Answer: Vendors are not required to run more than one workstream concurrently as part of the data developed for the proposal. The amount of data collected is up to the Offeror. There must be enough data concerning job interactions to determine what to bid to meet the proposed throughput performance as well as produce a convincing projection methodology as part of the RFP response.

Section J.1.4.2.1, General Comments states:

"In the ideal case, throughput benchmark measurements are taken on the systems proposed for delivery using the same queuing and scheduling software being proposed for the installed system. It is understood that realization of the ideal case is highly unlikely. Thus, it is generally expected that Contractors will take performance measurements on systems with the software scheduling and queuing infrastructures

currently available. After checking for interactions where the proposed R&D HPCS implies shared IT architecture components between workstream suites, projection methodologies may be used to produce the proposed configuration."

Thus, it is recognized that Offerors may not have the hardware and/or personnel resources to run the full throughput suite in preparing the RFP response. Still, the Government believes that in general, more data can be more convincing. Offerors must decide the balance between the effort to produce more data and the resources they have to produce it.

Question 244. An attempt to run the 52 day simulation as part of the throughput test for

WS3(HIMF), the run ends at 19.3 days. Is there any other criteria (convergence, for example) which can cause this premature ending of this run?

Answer: The simulation should run for the number of model time units specified in the ocean_solo_nml namelist in input.nml. Premature end is an error condition.

WS3 has been run successfully for the throughput simulation length on the NOAA SGI Altix platform using ifort.8.1.018 and 180 processors. See the ascii output for the throughput run in output.ws123.rfp.tgz on the rdhpcs website following the links:

Benchmark Information -> Complete scaling and throughput ascii output for workstreams 1, 2 and 3

Note that the RFP release of the source code for workstreams 1, 2 and 3 updates:

HIM/ocean_drivers/HIM_surface_forcing.F90

The original version contains a sign error at line 528. The erroneous sign will cause the model to become unstable.

Offerors are strongly advised to review the README files accompanying the RFP benchmark source release.

Question 245. 1. We interpret the answer to question 166 to say that if a single resource is proposed to meet the needs of (for example) WS1, 2, and 3 that for the purpose of both the SLT calculation and the Throughput Benchmark the combined resource **must** be proportioned according to the funding profile in C.4.3 Table 1 - that is exactly 4/14ths of the resource for WS1, 6/14ths for WS2, and 4/14ths for WS3. Is this correct?

If it is correct may we assume that the resource can be divided either spatially or temporally for the SLT calculation (e.g. WS1 uses 4/14ths of the resource for all of the year or all of the resource for 4/14ths of the year)?

2. Now suppose that a workstream (e.g. WS1) runs so well that 16 instances run in the same time as the required 8 instances for the Throughput Benchmark when constrained to using the 4/14ths of the combined resource. And suppose further that applying the

unused part of that 4/14ths of the resource from WS1 to the 6/14ths of the resource intended for WS2 allows WS2 to run twice as fast, thereby resulting in a significant decrease in the total of the Throughput benchmark for WS1, 2 and 3 combined.

Are we to understand from the answer to question 166 that the government requires the combined resource to be proportioned as the ratio of funding streams even if the combined Throughput Benchmarks on the combined resource run faster if the resource is proportioned on a performance ratio rather than a funding ratio?

If this is not the intent, may the offerer proportion the machine in such a way so as to minimize the combined Throughput Benchmarks of 1, 2, and 3 rather than by the financial ratio in C 4.2 Table 1?

3. Must the Throughput Benchmark used in the SLT calculation (C.6.1.2) be run in

exactly the same manner as reported for the Throughput Benchmark (J.1.4.2) result, or for example may the SLT calculation be done assuming multiple instances of Throughput Benchmarks are run simultaneously (thus potentially yielding a higher SLT result)?

4. Which result will be weighted higher in evaluating an offering - the Throughput Benchmark or the SLT number - and in what ratio?

Answer: Answer to follow.

Question 246. Were there important changes to benchmark source between the code released for the RFI and the source released for the RFP?

Answer: Yes. It is essential that Offerors review the README files accompanying the release of the RFP benchmark source. For example, as released with the RFI, there was a sign error in one of the forcing terms for the HIMF benchmark at line 528 of:

HIM/ocean_drivers/HIM_surface_forcing.F90

This sign error will cause the model to become unstable at simulation lengths longer

than the scaling study (e.g. the 52 day, 18 hour simulation length constituting the throughput run for WS3).

The correction is present and documented in the RFP release of the benchmark source.

Question 247. Amendment 2 changed paragraph C.5.2.4 to allow the HSMS nearline tier to be either disk or tape. The HSMS benchmark instructions at J.1.4.4.1 still specifically refer to locating all files on tape. Will the HSMS benchmark change to reflect the possibility of using disk as a portion of the near-line tier?

Answer: Answer to follow.

Question 248. We have several questions relating to connectivity requirements:

Q1 - Regarding the data transfer requirements addressed in C.5.2.5, C.5.4.3, and C.5.4.4:

Please confirm that the government is responsible for providing the necessary WAN connectivity to transfer all data listed in Table IIa and Table V from the source (whether the OCCS, backup OCCS, or the CEMSCS computer) to a) the Skaggs facility in Boulder(WS7-9 only), b) the NASA IV&V facility at Fairmont, and c) any

other government provided facility offered for hosting the HPCS in the DC Metro area. Please further confirm that the contractor is responsible for WAN costs only in the case that contractor-provided facilities are used for hosting HPCS under this contract.

Q2 - Please confirm that the data referenced in Table IIa and the data referenced in Table V are different data, so that the total data to be handled by the sites hosting HPCS for

workstreams 4-6 and 7-9 is the sum of the figures in those tables.

Q3 - Please confirm that paragraph C.5.4.4 refers to the local interface within the facility at Fairmont or Gaithersburg between a contractor-provided HSMS supporting workstreams 4-5-6 located at one of those facilities and the respective Primary or Backup OCCS interface point within the facility.

Q4 - Please clarify who is responsible for the connection between NCEP Users at the NOAA Science Center and the workstream 4-5-6 HPCS. Is this NOAA or the contractor? Does the same apply when the NCEP users migrate to the College Park site? Who will be responsible for service during the time that users are at both sites during the transition?

Q5 - If NOAA is responsible for the NCEP User to WS4-6 HPCS data connection,

how are the responsibilities for user session performance delineated?

Answer: Q1 - A. The data to support workstreams 7-8-9 is received at DSRC from a variety of sources. The DSRC has a satellite dish to download GOES data directly. Other data is obtained directly from their sources. ETA and GFS data is received via the NWS telecommunications gateway. Data is received from the AWIPS operational feed. If the vendor chooses to site the computational and storage resources that support Workstreams 7-8-9 in the DSRC then the Government will continue to provide the communications.

B and C: Data must be supplied between the OCCS and HSMS and between the HSMS and the system running the pre-production code, typically the operational backup. NOAA will provide transport for all data written from or read to these systems between these systems and the DC MAN. The vendor will be responsible for transport between the HSMS supporting these functions and the DC MAN.

In the event the vendor chooses to locate this portion of the HSMS in Fairmont, NOAA will transport the OCCS run histories data from the operational system to the HSMS. NOAA provisions the link between the operational and backup at a level necessary to support the operational requirement with high reliability and low latency. Typically there is sufficient bandwidth available to also transport the pre-operational data stream (Backup read and write in Table IIa) that is not used to support core

operational functions. Over any given period of time this available capacity will fluctuate. In addition, over time as that data grows and incorporates increasing amounts of satellite data, it is expected that this available capacity will no longer be sufficient to support the pre-operational data stream. In the event that that the available capacity for this link is exceeded, the vendor will be responsible for transporting the pre-operational data stream from Fairmont to the DC MAN or the backup computer, at the vendor's discretion.

Satellite data provided by CEMSCS is available on the DC MAN. Level II radar data is also available on the DC MAN.

The contractor is responsible for all WAN costs associated with a vendor provided facility. The contractor shall be responsible for all WAN costs that exceed GFE regardless whether the connection is required to a Government or vendor provided facility.

Q2 - There is an overlap of data between Tables IIa and V. Do not use the sum of both tables to calculate data storage requirements. Approximately 50% of the data contained in Table V is model data that is already contained in Table IIa. The remaining 50% of data in Table V is additional observational data from satellite and radar platforms that has not been ingested in the operational model runs.

Q3 - Section C.5.4.4 will be revised in a future amendment. Refer to Q1 answer to B and C (second para) for additional information regarding the interface to the HSMS.

Q4 - NOAA will provide the connection between the DC MAN and users at the NOAA Science Center. NOAA will provide the connection between the DC MAN and users at the future NOAA facility in College Park, MD. NOAA will provide the connections from the DC MAN to both facilities during the transition period.

Additional information regarding NOAA WAN connectivity which may support workstreams 4-6 will be made available in an amendment to the RFP.

Q5 - NOAA recognizes that that there may be a mixture of vendor and Government provided components and services along the communications path used to support WS 4-6, or for any of the workstreams. When problems occur, the vendor and Government will work together to identify the source(s) of any problem. Once the source of a problem has been identified, to the satisfaction of the Government, the party supplying that service or component will take responsibility for correcting the problem.

Question 249. Paragraph C.5.1.2 states that "The interactive resource should represent 5% to 15% of the LSC component." Please clarify whether that intent of this statement is that a) the IC resource is a subset of the LSC resource and therefore counts as resource that is used to deliver the promised SLTs, or b) the IC resource is

over and above the LSC resources and does not count toward delivering the promised SLTs.

Answer: Using IC resources to address system life throughput is addressed in #162.

The "5% to 15%" range is provided as guidance based upon NOAA experience.

Offerors can suggest other arrangements in their proposals.

Question 250. How long after submission of a question should we expect it to be posted on the website? Is there a requirement that only the primary point of contact post questions? Do we need to resubmit questions originally submitted on Monday?

Answer: The questions are posted after review by the Government. This is a manual process. We endeavor to post questions as quickly as possible.

There is no requirement that only the primary point of contact submit questions.

We will get the questions submitted Monday posted by 3/17.

Question 251. Amendment 2, C.11 and the recent answer to Question 112 indicate that the FSL subsystem HSMS configuration is being maintained by the current vendor for the first year of the new contract (i.e., through 9/6/2006). Please provide

answers to the following questions:

- a. Does the Government intend to continue to add data to the AML/J during FY 2006? If so, please provide an estimate of the data volume that will be stored in the AML/J on 9/6/2006. If not, please provide an estimate of the data volume that will be stored in the AML/J on 10/01/05.
- b. During FY2006, does the Government intend to operate both the legacy and new LSC elements in parallel?
- c. Assuming the contractor proposes a new HSMS solution for the FSL subsystem, please clarify whether or not the contractor has any requirement to provide access from the legacy LSC to the new HSMS during FY2006.
- d. Assuming the contractor proposes a new HSMS solution for the FSL subsystem, please clarify whether or not the contractor has any requirement to provide access from the new LSC to the AML/J during FY2006.
- e. Assuming that a contractor proposes to migrate data from the AML/J archive to a new HSMS solution, when can that migration begin (at the beginning of the new contract, or when the AML/J is available as GFE at the end of FY2006)?

f. Per Table V, section C.5.4.3 of the RFP (High bandwidth connectivity to model and observation data), the FSL workstreams require access to approximately 1 TB/day of model and observation data. If the Government intends to operate both the legacy and new LSC elements in parallel in FY2006, does it intend to provide separate feeds of these model and observation data to the two systems?

Answer: Answer to follow.

Question 252. In the answer to question 112, the government stated that the PRTN site will have installed a total of slightly more than 7PB of data in 5 StorageTek silos using 9840, 9940 and Titanium media. Are all 5 of the StorageTek silos, tape drives and media to be provided as GFE with regards to the RDHPCS acquisition. Also, can the government please provide us with the number of Titanium drives, and the number of Titanium media that will be in place at that time.

Answer: Answer to follow.

Question 253. When will the Goddard facilities be available for use under the contract?

Answer: Answer to follow.

Question 254. When and how will diagrams of the Goddard space be made available? How much raised floor space will be available, and how is it configured? How much non-raised-floor space will be available, and how is it configured?

Answer: Answer to follow.

Question 255. If build-out of additional resources is required, what is the process for funding this, how will it be executed, and what is the lead time required for its completion?

Answer: Answer to follow.

Question 256. Is the Government declaring that the Largo and Fairmont facilities described in the RFP are no longer available as GFE?

Answer: A decision to withdraw these facilities has not been made at this time. If such a decision is made, it will be announced immediately.

Question 257. When does current GFDL Contract with the incumbent Integrator expire? At the contract expiration date, will GFDL maintain title to the SGI Altix

systems or does the current Contractor retain title to the equipment?

Answer: Answer to follow.

Question 258. In developing the most cost effective solution and planning for a mid-life upgrade to the initial system delivery to GFDL, it is important for the bidders to understand the approximate disposition date for the SGI Altix hardware installed under the current contract. This information is necessary as bidders plan for the installation of the mid-life upgrade, given the current facilities environment at GFDL. Please provide the bidders with the anticipated removal date(s) for the SGI systems installed under the current contract.

Answer: Answer to follow.

Question 259. Will the existing Government and Contractor staff (under NASA contracts) at the Goddard facility be available to provide any support to the R&D HPCS systems housed there “operational monitoring, systems administration, system maintenance, or other?

Answer: Answer to follow.

Question 260. RE: Goddard Facility -

- 1) Who is the point of contact for questions concerning NISN?
- 2) How much GFE network bandwidth is NISN providing to support R&D HPCS data flows to and from
 - a) the MAX, b) Abilene, and c) the facility in Fairmont, WV?

Answer: Answer to follow.

Question 261. RE: Goddard Facility -

- 1) How much cooling will be available from existing air handlers at the time of occupancy?
- 2) Above and beyond the cooling available from existing air handlers, how much cooling is available from the existing infrastructure to support additional air handlers (how much additional chiller capacity is there)?

Answer: Answer to follow.

Question 262. RE: Goddard Facility -

- 1) Will the Offeror be responsible for the cost of electricity consumed by systems housed at Goddard? If so, how will this be calculated, and what is the current rate being paid for electricity by the facility?
- 2) How much power will be available at the time of occupancy from existing PDUs? How many receptacles (and what types) are available on existing PDUs?
- 3) Above and beyond power available from existing PDUs, how much power is available from the existing infrastructure to support additional PDUs?
- 4) How much UPS power will be available at the time of occupancy?
- 5) How much back-up generator power will be available at the time of occupancy?

Answer: Answer to follow.

Question 263. RE: PRTN Chillers - (1) What are the sizes and ages of the three chillers? (2) How does chilled water get from the mechanical room to the computer room? (3) Describe current responsibilities and schedules regarding maintenance of chillers. (4) How have chillers been acquired in the past?

Answer: (1) Section C.11.4.1 (Erroneously as C.10.9.1 in Amendment 2) indicates the following: Chiller #1 is 400 tons/ installed in 2000, Chiller #3 is 350 tons/ installed in 1996, and Chiller #4 is 225 tons/ installed in 1979; (2) The chilled water plumbing is indicated by the provided blueprints, including mechanical drawings from the 1979 (Computer) Facility construction (in which Chiller #4 was installed), the chiller upgrades in 1996 (chiller #3 installed) and 2000 (chiller #1 and new cooling towers installed), and the Raytheon mods in 2000 and 2002 (Note: The 75-kW generator installation for the chilled-water pump backup was included as part of the Boiler Plant Installation in 1998); (3) Daily operation and maintenance of the chillers is performed by Princeton University Maintenance staff, monthly preventative maintenance is provided under a government contract with York, and emergency response is provided by Princeton staff with fall back (with several hour response time) to York for more difficult problems; (4) by competitive Government contract

Question 264. RE: PRTN Power - (1) What is the voltage of the feed off of Route 1? (2) What is the size of the generator that was part of the infrastructure tour? (3)

How reliable is power off of Route 1?

Answer: Answer to follow.

Question 265. RE: PRTN - Describe the internet connection to Sayre Hall and to Internet 2.

Answer: There are two fiber cables between GFDL and Sayre Hall. One connects us to the Internet and is owned by Verizon. This cable is off-limits for any other uses. The other cable connects GFDL to the Princeton University network and is owned by Princeton. This cable has 9 single-mode pairs and 9 multi-mode pairs. Currently only 1 single-mode pair is in use. All fiber terminates in the computer room. There is conduit space between GFDL and Sayre Hall.

The fiber between GFDL and Sayre Hall connects two Cisco 3750 gigabit switches. We plan to move the GFDL end from the 3750 to a gigabit PIC module in the Juniper M7i.

Our I2 connection is via Princeton. Although GFDL has a gigabit connection to Sayre Hall, the connection from Sayre Hall to main campus is via a 100-Mbps microwave link. Princeton's I2 connection on main campus is also at 100 Mbps. The

connectivity to I2 is basically:

GFDL Juniper M7i to

GFDL Cisco 3750 to

Princeton Cisco 3750 to

Princeton microwave link to

Princeton Cisco 6513 to

Princeton Cisco 6506 to

MAGPI and Internet2

Question 266. RE: PRTN Backup for Cooling and Power - (1) What backups for the cooling facilities are in place? (2) What backups for power facilities are in place?

Answer: (1) None, other than historical N+1 capability; (2) Applicable UPS ride through, but no generator backup except to the chilled water pump.

Questions 267, 268, and 269. RE: Clarification Requested for Storage and Media at GFDL Site - There appear to be significant discrepancies in relation to the amount of storage that is to be installed, and available as GFE during the delivery schedule of the R & D systems. For instance, section C.11 Appendix C, Government Furnished Equipment (GFE) lists the Princeton site as having 4 StorageTek Powderhorn installed, and a total of 15,000 tape cartridges. Assuming that each of the SILOs can

hold approximately 5500 cartridges, is it to be assumed that there are 7,000 slots that are currently empty. The answer to question 112 implies that a fifth SILO is being installed, and that a total of 7PBs of data will reside in storage at the Princeton site at the start of FY 2007 that needs to be retained for the nine year life of the R & D contract. Why would a fifth silo be required, if it currently has 7,000 empty cartridge slots. Could the Government provide an accounting of total media by type, total number of tape drives by type, total number of slots in robotic tape systems, and total number of free slots anticipated to be in place at the start of FY 2007.

Answer: Answer to follow.

Questions 270, 271, and 272. Clarification of Data Generation Rates for Archiving –

a) Section C.5.2.6 provides a data generation profile for WS1 - WS3 of 5 TBs per day. The governments answer to question 36 expresses that over 10 TBs of data are written to its current tape archive system. What additional data is being written today that is not included in the data generation amount provided in section C.5.2.6. Is this additional data stream going to go away before the start of FY 2007, or is this data rate a new requirement that has not previously been accounted for in the RFP?

b) If this is a new storage requirement, does the additional amount of data generated scale

using the same equations as outlined under Section C.5.2.6?

c) Are there any other data storage requirements that have not previously been outlined in the RFP that need to be taken into account by vendors responding to the NOAA R & D RFP?

Answer: Answer to follow.

Question 273. RE: Clarification of Nearline Storage Needs - Amendment 2 changed Section C.5.2.4 to allow the HSMS near-line tier to be either disk or tape. Paragraph 9, sentence number two stated: -The Government desires that files that haven't been accessed in one year will be migrated from the nearline storage to the offline storage-. This statement implies that the nearline storage system is required to be sized such that all data generated within a rolling one year period is to be stored in nearline storage. If the nearline storage being proposed is disk, then using the baseline data generation profiled in section C.5.2.6 of 5 TBs of data generated per day would lead to the requirement that starting in FY 2007 the nearline storage (as stored on disk) would need to be a minimum of 1.825 PBs, and would need to grow according to the data generation growth equations provided in Section C.5.2.6. Is this assumption correct?

Answer: Answer to follow.

Question 274. RE: PRTN - Describe the Raytheon/SGI support functions that are currently in place at the Princeton facility.

Answer: As they relate to facilities, Raytheon/SGI maintains the UPSs. Raytheon maintains the air handlers for the first year after installation and then turns maintenance over to the building owner, Princeton University, who performs routine and periodic maintenance.

Question 275. RE: PRTN: Computer Room Layout - What will be the layout of the equipment in the Computer Room after the mid-contract upgrade?

Answer: See Figure 3a, which is available from the Contracting Officer upon request.

Question 276. RE: PRTN UPS Issues - (1) What is the ride-through capability for the UPSs for the current systems? (2) What about the shelving currently located in the UPS Room in the north corner of the Computer Room?

Answer: Answer to follow.

Question 277. RE: PRTN Air Handlers - (1) Have there been any air handler

failures recently? (2) Where is air handler #6?

Answer: (1) No, there have been no known air handler failures since the start of the Raytheon contract; (2) This is the small air handler located in the ceiling of the Printer Room; (3) No.

Question 278. RE: PRTN Workload as Operational - Does the computer run a production (operational) workload? Does it require N+1 backup capabilities?

Answer: The government executes long-running climate jobs at PRTN that could be viewed as production class jobs, but it does not run an operational workload in the sense of the National Weather Service's operations. N+1 for the chillers is the Government's requirement at PRTN.

Question 279. PRTN: Who owns the Brocades? Are they GFE?

Answer: They are GFE. The brocade switches are listed in the GFE listing in the RFP (C.12 Appendix C [C.11 in Amendment 2]) with Part No. designation "FC-SWITCH-16", description "16 port 1Gb FC switch (8Cu/8Optical)" and a quantity of 12 under "Hierarchical Storage Management System."

Question 280. PRTN: Is there an environmental report available on the existing building? Any asbestos materials?

Answer: The Government has contracted for a full asbestos survey of the GFDL Main Building; the report from this survey is still in draft form and is not available for public release. The report indicates that asbestos is present in locations throughout the Main Building. Regarding areas of the building that may be relevant to this solicitation, asbestos was abated this past year from room 141 (the Mechanical Room) and room 143, which contains the Main Building Substation and the phone system. In 1985, asbestos was abated from the region of the Main Building between the 3rd-floor ceiling and the roof. The Government has not conducted an asbestos study of the Computer Building, which was built in 1979.

Question 281. PRTN: Are there electric utility bills (for the past year) available for review? Other utility bills?

Answer: As indicated in Amendment 2, the Government will provide a spreadsheet containing a recent history of electric utility monthly costs and consumption upon request to the Contracting Officer.

Question 282. PRTN: Is there a visualization lab? Why was it not included on the tour?

Answer: Yes, there is a visualization lab. During the Q&A session at the PRTN site visit, we announced a tour of the lab, which was provided to interested parties at the conclusion of the meeting. Workstations in the viz lab are connected to an Onyx system in the Computer Room as part of the current contract; however, this specific functionality is not a requirement in the RFP.

Question 283. PRTN LAN Connectivity - What LAN connectivity is available for additional processors - 100 Mbps port? Gigabit ports?

Answer: Answer to follow.

Question 284. PRTN Airflow units - (a.) What are the model numbers of the airflow computer room A.C. units? (b.) Are there humidifiers and reheat coils in the airflow units?

Answer: (a.) Section C.12 (“C.11 Appendix C” in Amendment 2) gives the Manufacturer and Part Numbers as follows: the 35-ton CRAC 1,2,3,4, 5, and 7 are DataFlow Part No. CCT-60C4, the 35-ton CRAC 8 is APC Part No. CCT-60C4, and the 3-ton unit is APC Part No. CM-3.0-W-BC-D. (b.) Yes. However, PRTN normally uses the electric steam boiler for humidification.

Question 285. PRTN Drains / Water for Cooling Tower – (a.) Are there floor drains in the pipe trough under the raised floor? Sump pumps? There appear to be no water sensors in the pit. (b.) Is there a back-up system for make-up water to the cooling towers? If so, what is its source and what is its capacity?

Answer: (a.) There is a drain in the pipe trough, but no sump pump or water sensors. (b.) No. Historically PRTN used backup make-up water from a nearby canal, but this past year the use of that supply has become unreliable, probably due to the age of the underground pipes.

Question 290. PRTN UPS Issues - (a.) What is the present UPS load for both systems? (b.) Are the air handlers powered by the UPS? (c.) What is the UPS battery ride-through time? (d.) What is the vintage of the UPS? Vintage of the batteries? (e.) What is the type of the UPS battery: VRLA, GEL, etc?

Answer: (a.) Currently, the loads on the 500- and 225-kVA UPSs are 386 and 109 kVA respectively, for a total load of 495 kVA. It is projected that the UPSs will be operating close to full (80%) capacity once the April 2005 upgrade is completed. (b.) No. (c.) As Section C.11.7.2 (Referenced as C.10.12.2 in Amendment 2) indicates, the ride-through times for the 500- and 225-kVA UPSs currently are shown on the read-out panels to be 16 and 25 minutes, respectively. Replacing the old batteries in

the 500-kVA UPS [reference (d.) below] is expected to increase the ride-through time significantly for this UPS. (d.) The 500-kVA UPS was installed with the Cray T932 in 1997. The 225-kVA UPS was installed with the Cray T94 in 1998. In the 500-kVA UPS, three of the four banks of batteries are brand new, and the fourth bank is over four years old; Raytheon is planning to replace the latter bank during the next two weeks. The 225 KVA UPS batteries are six months old. (e.) Wet cell.

Question 292. Section C.10.1.3 describes the GFDL LSC as primarily composed of SGI Origin processors. However, it appears that the baseline performance is computed on the soon to be installed SGI Altix systems. Please clarify that the government desires a substantial performance increase (section C.5.1.1) over the soon to be installed Altix systems. If it is based on the Altix systems, can the government please describe the size of each Altix system that will be installed just as they have described the origin systems (node size, memory, etc)?

Answer: The desired "substantial performance increase" for all workstreams is with respect to the baseline throughput values as provided by the RFP. Specifically for workstreams 1, 2 and 3, the performance increment is with respect to the performance baseline developed using an SGI Altix system. The configuration used to develop the baseline throughput is clearly defined at the top of Section J.1.4.2.4.

Given the timeline for finalizing the upgrade to the Princeton facilities relative to that

required for the RFP response, the Government does not anticipate providing further information concerning the Altix upgrade. Moreover, the details that upgrade will not change the throughput baseline for workstreams 1, 2 or 3. It is against the RFP throughput baseline that proposed performance will be evaluated.

Offerors are reminded that the goal of the RFP is to obtain the greatest amount of computational throughput for the dollars available to a workstream in the context of the workstream's archive and other "non-computational" requirements. The Government has purposely avoided specifying minimum performance increments to allow the Offeror the greatest flexibility to optimize the proposed configuration within the funding profile

Question 293. PRTN Electrical Issues / Scheduled Maintenance - (a.) Are there TVSS (Transient Voltage Surge Suppressor) devices on the existing switchboards? (b.) Has a short circuit/protective device coordination study been performed? (c.) Is the raised floor grounded? If so, at what interval? (d.) What is the schedule/interval of maintenance on the following equipment: UPS, switchboards, transformers, chillers, and air handlers?

Answer: (a.) No. (b.) No. (c.) Yes, the raised floor uses a bolted stringer construction that is grounded with a single ought (Zero) grounded loop. (d.) UPS: Raytheon contract for twice annual preventative maintenance plus on-call service; Switchgear:

Cleaning of breakers when necessary; Transformers: No explicit program for preventative maintenance; Chillers: Monthly preventative maintenance is provided under a government contract with York along with annual servicing as required and on-call support; Air Handlers: Air handlers remain under warranty for one year from the date of purchase. Princeton University performs routine maintenance and service on them once they are no longer covered under warranty

Question 294. PRTN: Loading and Circuits on PDUs - What is the current loading and number of circuits on each of the PDUs and what is the projected April 2005 loading?

Answer: The three (3) 225-kVA United Power PDUs each have four breaker panels, where each panel has 42 circuits. The Liebert 125-kVA PDU has 3 breaker panels, each with 42 circuits. The EPE 125-kVA PDU has 2 breaker panels, each with 42 circuits. After April 2005, all of the PDUs will be 100% full.

Question 295. PRTN NLR Connection / High-Speed Network Access - (a.) Will a National Lambda Rail (NLR) connection be available? When? Will it terminate in the Computer Room? (b.) Is any other high-speed network available in the area? What is the approximate distance from the GFDL site?

Answer: Answer to follow.

Question 296. PRTN GFDL Power Available /UPS Ride-Through - (a.) What is GFDL's share of the total power feed from PSE&G (out of the 2.5 MVA)? (b.) How long can the UPSs support the current system in the case of power failure?

Answer: (a.) As Section C.11.3.1 (C.10.8.1 in Amendment 2) indicates, Princeton has indicated that GFDL's share is 1430 kW due to an observed peak demand of 1070 kW for other users of the 2.5 MVA substation. (b.) Please see the response to Question # 290.

Question 297. RTN UPS Issues – (a) What is the age of the Raytheon UPS? (b.) Are the air handlers on the UPS? (c.) What is the ride-through time on the current UPSs?

Answer: (a.) Please see the answer to Question #290. (b.) No. (c.) Please see the answer to Question #290.

Question 298. PRTN Equipment DeInstalled After FY2006 - On the April 2005 layout, please note the Raytheon equipment that will be de-installed at the end of FY2006.

Answer: Answer to follow.

Question 299. Paragraph C.5.6.1 in Amendment 2 states a requirement that the Contractor maintain any UPS equipment used to support systems under this contract. In order to obtain quotes from the manufacturer of the existing UPS units at PRTN, the manufacturer is requiring us to provide serial numbers. Please provide the serial numbers for all of the relevant UPS units at PRTN so that we can obtain the necessary quotes.

Answer: PRTN: The MGE 500-kVA UPS has the following identifiers: Model No. '72-130104-00 EPS 6500/44,66' and Serial No. 69937-01; the MGE 225-kVA UPS has the following identifiers: Model No. '72-130101-01 EPS 6225/44,66' and Serial No. 200834-01.

Question 300. GRBLT Nearby GigaPops / Telco Demark - 1.) Which network providers have nearby GigaPops? Please provide name of provider and the distance.

2.) Where is the telco demark and how close is it to the NOAA-provided space?

3.) Is there sm Fiber available to extend the telco demark to the provided space? What will be the cost to the Offeror to do this?

Answer: 1.) University of Maryland College Park/~10miles

2.) Bldg. 1/~1 mile

3.) Yes, \$1,500/per pair of strands for standard single mode. A higher quality connection is required to support some high-bandwidth WDM and that may cause the price to be higher.

Question 308. GRBLT Fire Suppression - What type of fire suppression systems are/will be available in the facility?

Answer: Wet pipe sprinklers in room, fire detection above ceiling, in room and below raised floor.

Question 311. GRBLT Age of building - How old is the building containing the provided space?

Answer: This portion of the building was completed in 1/1996.

Question 313. GRBLT Walls Exterior to Room - During the planned room reconstruction, would the outside walls surrounding the offered space be run down to the concrete floor below the raised floor?

Answer: Perimeter walls will go slab to slab.

Question 316. GRBLT Maximum power - What is the maximum power available to the Offeror?

Answer: 1200 KVA on the 480 V system and another 1200 KVA on the 208V system.

Question 317. GRBLT Rating/size of electric utility - What is the rating/size of the electric utility service?

Answer: Service entrance equipment consists of 2 double-ended metal-clad indoor unit substations, one rated 1500KVA at 480V 3 phase 4 wire, and one 1500KVA at 208V 3 phase 4 wire.

Question 318. GRBLT Redundancy - Are there redundant utility AC feeds?

Answer: Not within this space.

Question 321. GRBLT Shutdowns for maintenance - Does NASA schedule shutdowns of its utilities in order to perform maintenance and, if so, how often?

Answer: There are scheduled outages for utility services. Generally these do not cause downtime for the users.

Question 322. GRBLT Cooling Capacity - Regarding the 130 tons of cooling capacity that is indicated to be available to the Offeror, is this available as chilled water or condenser water?

Answer: Chilled Water

Question 325. GRBLT Power service - What is the power service to the building? What will be the power service to the offered space?

Answer: Power will be one 800 amp main distribution panel rated at 480 V, 3 phase, 3 wire.

Question 326. GRBLT Floor - When NASA turns the space over to NOAA, will the space be demised with a slab-to-slab partition, including a barrier below the raised floor?

Answer: Yes

Question 327. GRBLT Electrical distribution panel - Will NASA terminate their work with an electrical distribution panel serving solely the NOAA space?

Answer: Yes

Question 330. GRBLT Breakers Config. - The available 520 kVA of power is assumed to be provided with 480 volt single or multiple breakers in the existing space. Is this correct?

Answer: Service entrance equipment consists of 2 double-ended metal-clad indoor unit substations, one rated 1500KVA at 480V 3 phase 4 wire, and one 1500KVA at 208V 3 phase 4 wire. Power will be one 800 amp main distribution panel rated at 480 V, 3 phase, 3 wire.

Question 338. GRBLT Networking - 1.) Who is the point of contact for questions concerning NISN? 2.) How much GFE network bandwidth is NISN providing to support RDHPCS data flows to and from the following locations: The MAX, Abilene, The facility in Fairmont, WV?

Answer: GSFC POC is Mary Ellen Shoe 301-286-8920. NISN provides wide area network services only.

The MAX - A) NISN has 1Gbps dedicated connection for general users at GSFC. B) 2Gbps with 1Gbps backup (internal R&D effort by various projects). GSFC connection to Abilene is via the MAX at University of Maryland College Park.

Abilene - A) NISN has 1Gbps dedicated connection for general users at GSFC. B) 2Gbps with 1Gbps backup (internal R&D effort by various projects). GSFC connection to Abilene is via the MAX at University of Maryland College Park.

The facility in Fairmont, WV - Currently 2MB (T1 + fractional). In the process of upgrading to a DS-3. OC-12 circuit is part of their research network and is only through West Virginia University and does not connect through NISN.

Question 341. Fairmont and Largo - Is the Government declaring that the Largo and Fairmont facilities described in the RFP are no longer available as GFE facilities?

Answer: A decision to withdraw these facilities has not been made at this time. If such a decision is made, it will be announced immediately.

Question 343. GRBLT Wall Exterior to Room - What are the plans for isolating the room(s) from other surrounding rooms?

Answer: The room is to be reconfigured to a single space that is separated by walls that run slab to slab.

Question 344. GRBLT Power / Cooling to Room - How will power and cooling be brought to the Offered space?

Answer: Power brought to main distribution panel in space. Chilled water piped below the raised floor into space and capped.

Question 347 GRBLT Possible Building Code Issue - I note that the room in the upper right corner of the provided space appears to contain a possible building code violation [I did not understand the issue, although it could relate to single entry to the room!]. Will this be resolved?

Answer: The space will be configured as one room. A safety review will be done to establish specific exit requirements.

Question 348. GRBLT Cost of Network Connection - Who will be responsible to pay for the cost of the network connection to the provided space - the Government or the Contractor?

Answer: Contractor

Question 350. GRBLT Costing Policy for Dedicated Bandwidth - From my dealings with NASA regarding other projects at GSFC, I have found that, whereas limited internet access without guarantees for availability is essentially free, but internet access with guaranteed allocated bandwidth requirements can be quite expensive. Please explain your policies regarding this. How can Offerors develop cost proposals in response to this procurement so as to accurately project costs to meet the bandwidth requirements that they consider to be necessary to meet the Government's requirements? In other words, how do we, as Offerors, minimize our cost risk regarding networking requirements in preparing our proposals?

Answer: Campus connectivity without guarantees is available at a modest cost per connection, if the requirements exceed capacity then special LAN/WAN are developed by the requesting customer(s) and they absorb the full cost of services. General LAN/WAN is provided by the Centerwide Network Environment (CNE) Project managed by Curt Suprock 301-286-6196. His office is willing to provide responses to questions of network services. If the questions are general in nature then Curt Suprock is the primary contact. If the questions are engineering in nature then Dave Smith 301-286-0474 is the primary.

Question 351. GRBLT NISN Point of Contact - Is there a NISN point of contact for better information regarding networking issues?

Answer: GSFC POC is Mary Ellen Shoe 301-286-8920. NISN provides wide area network services only.

Question 355. GRBLT Switchgear Age and Configuration - What is the configuration and age of switchgear in the building that will service the provided space?

Answer: Service entrance equipment consists of 2 double-ended metal-clad indoor unit substations, one rated 1500KVA at 480V 3 phase 4 wire, and one 1500KVA at 208V 3 phase 4 wire. The equipment was installed in 1996 with the building construction.

Question 356. GRBLT Chilled Water Plant / Chilled Water Pipe - Where is the chilled water plant located that services the building? What is the capacity of the pipe that feeds chilled water to the building?

Answer: The Central plant, building 24, is to the west of building 28; a 12" pipe.

Question 357. Clarification of WS1-3 data sources - Section C.5.4.3 discusses connectivity to model and observation data and states that ...If the workstreams are targeted for computational platforms not co-located with their current data source bidders need to provide the required bandwidth. This section describes the data sources and data ingest rates for ws4-6 and ws7-9. However, WS1, 2 and 3 are not discussed. Where do ws1-3 obtain their data from and at what rates?

Answer: While workstreams 1-3 may import data gathered through observations for model parameterizations (such as topography data), the models do not operate on observational data (i.e. there is no data "ingest" such as that associated with workstreams 4-9).

The INPUT directory supplied with workstreams 1 and 2 provide an idea of how much data is required to run the next segment of the model. To minimize startup costs, workstream 3 generates an idealized set of values internally. Actual data sizes required to restart the workstream 3 model are on the order of the workstream 2, CM2-HR model.

Restart and model parameter files are all located in the archive associated with workstreams 1, 2 and 3.

Question 358. Use of Scaling Study Data - Section M.3.1 states items used to evaluate performance may include the results of the benchmark scaling study. However, in section J.1.4.1, subitem ii, Scaling Study, states that the component scaling is not evaluated in and of itself. Can the Government please clarify if the scaling study will be used to evaluate performance or will it be used for informational purposes only?

Answer: The Government believes the statements in Section J.1.4.1 and M.3.1 to be consistent. Scaling performance in and of itself is not an evaluation criterion. The scaling data will be used as stated in Section J.1.4.1:

"The purpose of this test is to aid the understanding of performance projections and the intended model performance point for the offered system."

The quality of the performance projections, as well as the model performance point for the offered system, are parts of the evaluation as per Section M

Question 360. PRTN: (a.) Why is there a 5-degree grade on the hardpan floors in the Computer Room? (b.) How is potential water damage prevented/detected? (c.) Why does the security camera system only retain video for one week and not for three months?

Answer: (a.) The hard pan is level, except at the edge of the computer room floor, where it has a slight dip, probably to accommodate possible runoff. (b.) Drains in trough. Water

detection sensors are available, but have not been installed. (c.) See Answer to Question #264 regarding reliability of power.

Question 361. PRTN: (a.) Is your alarm integrated into the fire suppression system? (b.) Why does the security camera system only retain video for one week and not for three months?

Answer: (a.) Yes. (b.) The security camera system retains video for 31 days because that is the Government's requirement at this site.

Question 362. Section C.5.1.1 states: "The contractor shall provide a LSC component at a substantial increase in sustained throughput over NOAA's current supercomputers described in Appendix A." It is understood that the government did not run all the copies required for each workstream. It is also understood that other codes may run on these systems. However, the provided benchmark codes are the ones that bidders must use to evaluate the throughput, and determine the performance improvement that can be provided.

For the WS4-6 system, there are enough processors in their current cluster so that all the jobs can run concurrently. Each job runs on separate processors, but to obtain an accurate estimate of throughput, a bidder would need to account for the effect of contention for shared resources, such as the file-system. A bidder can then use the computed throughput to compare against proposed throughput, and ensure that the bidder is providing a performance improvement.

However, for GFDL WS1-3, this kind of analysis is problematic. The answer to question 172 states: "The April 2005 configuration for the Princeton site LSC will be comprised of 512, 256 and 96 processor SGI Altix Systems. Thus, the configuration is a cluster of systems and there is no single system image. Further, the Government currently has no plans to run applications outside a single system's numalink communication fabric."

Given the information provided, it is not feasible to develop a realistic baseline throughput number, using the throughput benchmark jobs. For example, a bidder does not know how many of the nodes have 96 processors, how many have 256 processors, and how many have 512 processors. Furthermore, the throughput job for WS1 uses 135 processors, the throughput job for WS2 uses 240 processors, and the throughput job for WS3 uses 180 processors. Thus none of these jobs can run on 96-processor nodes. On the 256-processor nodes, a single WS2 job can run, but 16 processors are unused; or a single WS1 job can run and 121 processors are unused; or a single work WS3 job can run and 86 processors are unused. The same logic can be applied to the 512-processor nodes.

Thus it is not feasible to compute a baseline throughput for the current GFDL system, which is required as a reference point to evaluate the performance improvement in a bidders proposal.

- a) Given the above analysis, can the Government please provide guidance on how a bidder can compute a baseline number for WS1-3, such that bidders can be responsive to Section C.5.1.1, to provide a substantial improvement?
- b) Perhaps, if the Government considered a redefinition of the processor counts and reran

a single instance of the throughput jobs for workstreams 1-3 so that they can all run concurrently on the 3296 processor Altix cluster, a more suitable baseline could be developed.

Answer: The Government disagrees with the analysis. The baseline, whether "real" or "synthetic", has been "computed" in the RFP. For the purposes of the RFP:

1) The throughput baseline for workstream 1 is 8 instances of CM2-ESM with a simulation time of 420 days and a throughput wallclock time of 21825 seconds.

2) The throughput baseline for workstream 2 is 6 instances of CM2-HR with a simulation time of 51 days and a throughput wallclock time of 21001 seconds

3) The throughput baseline for workstream 3 is 4 instances of HIM-VHR with a simulation time of 52 days, 18 hours and a throughput wallclock time of 21680 seconds

There is nothing further the Offeror needs to calculate with respect to the baseline. The baseline is provided.

The goal of the RFP is to obtain the greatest amount of computational throughput for the dollars available to a workstream in the context of the workstream's archive and other "non-computational" requirements. The Government has purposely avoided specifying minimum performance increments to allow the Offeror the greatest flexibility to optimize the proposed configuration within the funding profile.

The PE count performance information provided with the RFP is merely to give some idea of model performance on a given technology at a given point in time. The Government fails to understand the relevance of PE counts used for current performance information to processors and architectures to be available years in the future.

The Government has no plans to redefine the benchmark baseline for workstreams 1, 2 or 3 at this time.